FREE FIELD GUIDE · REVENUE EXPERTS

AI visibility: is your company the answer when buyers ask AI?

A field guide to seeing what ChatGPT, Claude and Perplexity say about you, deciding the right fix, and becoming the answer. Five prompts, and the judgment to use them.

Short answer. AI visibility is whether AI assistants name your company when buyers ask who to buy from. To check yours, ask ChatGPT, Claude and Perplexity the questions your buyers really ask, three times each, and record who gets named and cited. Most B2B companies are named in only a few, or none. The five prompts below do exactly that, then help you pick the one fix that moves the number: rewrite a page, build a new one, write an article, or get onto the lists AI trusts. This is answer engine optimization (AEO), run by hand.

What AI visibility means, and why it matters now

Your buyers used to open Google, scan ten links, and click one. Now many open ChatGPT, Claude or Perplexity, ask a question in plain words, and get back a short answer that names two or three companies. If you are not one of them, you were never on the list, and there is no second page to be found on. The job has changed from ranking to being named. This guide shows you whether the assistants name you today, why they name whoever they name, and the exact moves to become the answer. You can run all of it yourself, for free, this week.

How does an AI assistant decide who to recommend?

Three forces decide it, and every prompt below targets one. Hold the whole picture first.

- **It quotes the page that answers the exact question**, directly and early, with specifics it can lift. A broad "our services" page almost never qualifies; a page built around one question does. This is why the fix is so often a page, not a slogan.
- **It leans on other people's lists far more than on your own site.** Directories, best-of lists, review sites and comparison pages carry most of the weight. In our field service study

(October 2026), 1,181 of 1,493 citations went to sites other than the vendors themselves. Sometimes a seat on one list beats anything you could publish.

- **Being named and being cited are two different things.** An assistant can describe you without linking you, or quote your page without saying your name. In our NRS study (September 2026), 14 answers used a company's own pages and never named the company once. You will measure both, because fixing them is different work.

What is the five-step loop?

See who gets named. Diagnose why. Decide the right fix. Write it or get listed. Re-check next month. One prompt per step, plus the one decision most people get wrong. Do it once by hand and you will understand AI visibility better than most agencies selling it.

Step 1. Which 20 questions should you try to win?

What you get: 20 questions your buyers ask an assistant before they choose, in their words, grouped by how close they are to buying. **Why it matters:** you can only win questions you know buyers ask, and your company name is never one of them; this list is the measuring stick for everything after it. Run this once in any assistant.

```
I sell [WHAT YOU DO] to [WHO BUYS IT]. Think about how buyers in this category
search before they choose a vendor, and give me 20 questions they would ask an
AI assistant, in their own words, WITHOUT my company name in them.

Mix these kinds: "best [category] for [situation]", "how do I fix [problem]",
"[option A] vs [option B]", "what does [thing] cost", and "is [approach] worth
it". Group them by buying stage: just looking, comparing, ready to choose. Keep
each to one line.
```

Do this with it: keep the 20. They are your fixed test set; they do not change unless your market does.

Step 2. Does AI name you today?

What you get: for each question, who the assistant named and in what order, which sources it actually opened, whether it named you, and a line on why the winners were easy to pick. **Why it matters:** this is your baseline, the single number that tells you whether you have a problem. Run it one question at a time, in ChatGPT, then Claude, then Perplexity, three times each in a fresh chat, and keep score yourself.

```
Turn on web search and answer this question exactly as you would for a buyer
asking it cold: "[ONE QUESTION FROM YOUR LIST]".

After your answer, give me this short report:
1. The companies you named, in the order you named them.
2. The sources you actually opened (links only; if you did not search, say so).
3. Did you name [COMPANY] ([WEBSITE])? If not, which company got the spot I want?
4. In one line each, why those companies were easy to recommend here: what on
   their own pages, or on other sites, made them the obvious answer.

Do not guess to fill a gap. If you are unsure, write "unsure".
```

Do this with it: fill the scoresheet near the end. Your count out of 20, dated, is where you start. Everything else is about moving that number.

Step 3. Why does AI pick them and not you?

What you get: the three to five things the winning pages have that yours does not, the single most likely reason you lose, and a checklist, most important first. **Why it matters:** "we are invisible" is not actionable; "our page buries the answer in paragraph four and names no specifics" is. It feeds the one real decision, next.

```
Here is a question I want to be the answer to: "[QUESTION]".
Here is my page on this topic today: [PASTE THE PAGE TEXT, OR THE URL IF THIS
ASSISTANT CAN BROWSE].

Turn on web search, read the pages that currently get recommended for this
question, and tell me:
1. The 3 to 5 things those pages have that mine does not (for example: the answer
   in the first lines, a comparison table, specific dated numbers, a named
   author, mentions on other sites).
2. The single most likely reason an assistant picks them over me.
3. A checklist of what to change, most important first.

Use only what you can actually read; flag anything you cannot open. Do not write
anything yet.
```

New page, rewrite, or article: how do you choose?

You now have a gap and a checklist. Before you write a word, decide the right move, because the wrong one wastes weeks. Ask three questions in order.

- **Do you already have a page that is read, or ranks, for this topic?** If yes, improve that page. Do not publish a second page that competes with it; two weak pages on the same question beat neither. A rewrite usually wins faster than something new.
- **Is the question its own distinct thing, with nothing of yours that answers it directly?** Then write one new page, built around that question, answer first. This is why the fix is so often a new page: assistants quote the page that answers the exact question, and your homepage or services page answers twenty questions at once, so it answers none well enough to be quoted.
- **Did Step 2 show the assistant only ever cites third-party lists and comparisons for this question, almost never a vendor's own page?** Then writing anything on your own site will not win it. Your move is to get onto those lists (Step 5), not to write a page. It is the cheapest win, and the one most people miss.

So where does an article, or a blog post, fit? An article is the amplifier, not the answer. Use it for a timely angle, an opinion, or a study that earns attention and links, and point it at your answer page; it rarely becomes the evergreen answer itself, because it reads as news, not as the definitive page on the question. Use an article to attract and to earn mentions on other sites; use a page to be the answer. If you only have time for one, build the page, then write the article later to send links and readers to it.

Step 4. How do you write a page that wins the answer?

What you get: a publish-ready page (or the rewrite you chose), structured the way assistants quote. **Why answer-first:** the assistant lifts your opening lines; bury the answer and it quotes whoever put theirs on top. This is the one task a chatbot does genuinely well, so let it draft, then you check every fact and publish it yourself.

```
Write a publish-ready page that could become the AI answer to: "[QUESTION]".

Shape: the heading is the question in the buyer's words; the first 40 words are
the direct answer, before any wind-up; then the reasons, the steps, one
```

```
comparison table, and a three-question FAQ in plain text.

Use ONLY these facts, and mark anything you cannot verify as [unverified] rather
than inventing it:
- [REAL FACT]
- [REAL FACT]
- [A REAL NUMBER OR STUDY]

Do not add prices, guarantees, invented statistics or reviews, and do not call
any competitor bad. At the end, give me a title under 60 characters and a meta
description under 155.

This is a draft for me to review and publish myself.
```

Step 5. How do you get onto the sites AI trusts?

What you get: which third-party sites to target, whether they name you or a rival today, the realistic way onto each, and three ready outreach notes. **Why it matters:** go back to force two. For a large share of questions the assistant cites lists and directories, not vendor sites, so a seat on the right list is worth more than a month of writing. Paste the sites you saw cited in Step 2.

```
Here are the websites AI assistants cited when I ran my questions, with how
often each came up:
[PASTE THE SITES AND THEIR COUNTS THAT YOU COLLECTED IN STEP 2]

Ignore my competitors' own sites. For the best-of lists, directories, review
sites and trade pages, tell me for each one: does it name [COMPANY]; who it names
instead; and the realistic way to get included (an editorial pitch, a form, a
paid listing, a customer review).

Then write one short note, under 120 words, for the top three, each mentioning
one real thing on that page and making one clear ask. Use only contact details
the page shows publicly; flag any with none. These are drafts for me to send.
```

How do you keep score?

A chatbot cannot run all 20 questions at once, remember between chats, or save a file, so you keep the tally. Count how many of your 20 questions named you in each assistant. Write the number down and date it, then re-run Step 2 in a month and watch the number, not your feeling.

```
Question                      | ChatGPT names you? | Claude? | Perplexity? | Named instead
1. ............................ |       y / n       |  y / n  |    y / n    | ...........
2. ............................ |       y / n       |  y / n  |    y / n    | ...........
(repeat for all 20)           |                   |         |             |
TOTAL named of 20             |       __ / 20     | __ / 20 |   __ / 20   |
```

What will you have at the end?

A dated baseline of how many of your 20 questions name you in each assistant, a prioritized fix list, one page live that is built to be the answer, and a short path onto the sites the assistants trust. That is a real AI visibility program, run by hand, for the cost of your time. For the weekly news that changes these answers, keep an eye on our newsroom; for a check a week on your own site, the Lab walks you through one at a time.

What are the limits of a DIY check?

- AI answers change between runs and between a signed-in and a signed-out session, so run each key question three times and keep all three.
- Assistants name a company without linking it, and cite a page without naming the company, so a by-hand check shows you the shape, not an exact number.
- No prompt and no page can guarantee a citation. What moves the number is being readable to the crawlers, original evidence, clear sourcing, and answering the exact question first.

AI visibility FAQ

What is AI visibility?

AI visibility is how often AI assistants like ChatGPT, Claude and Perplexity name and cite your company when buyers ask about your category. It is the AI-answer version of showing up in search, and it is measured by asking real buyer questions and recording who gets named.

How do I check if ChatGPT recommends my company?

Ask ChatGPT, Claude and Perplexity your buyers' real questions, without your company name, three times each with web search on, and record whether you were named and which sources were cited. Prompt 2 above does this; run it for all 20 of your questions and tally the score.

Is AEO different from SEO?

They overlap. SEO is about ranking a page in a list of links; answer engine optimization (AEO) is about being the company an assistant names in its written answer. Good SEO still helps, because assistants read the live web, but AEO adds answer-first pages, clear sourcing, and mentions on the sites AI trusts.

Do I need a new page, or can I update an existing one?

If you already have a page that is read or ranks for the topic, update it rather than add a competing one. Write a new page only when the question is distinct and nothing of yours answers it directly. When only third-party lists get cited, get listed instead of writing anything.

How much does an AI visibility audit cost?

Our AI citation audit is $1,495: 50 buyer questions across ChatGPT, Claude and Perplexity, three runs each, 450 saved answers, every citation checked, and a fix list with owners, in 2 to 5 days. The DIY version in this guide is free.

When you want it measured, not estimated

This guide gives you a snapshot you run by hand, and the judgment to act on it. When you want the measured version, done for you, that is our AI citation audit: 50 buyer questions across ChatGPT, Claude and Perplexity, three runs each, 450 answers saved and dated, every citation opened and checked, and a fix list that already tells you, question by question, whether to rewrite, build a new page, or get listed, and who owns each fix. 2 to 5 days. $1,495. You keep the saved answers, so the number holds when someone asks where it came from.

Ready to see yours measured? Book a 20-minute call with John Bush, or tell us about your market on our contact page, and we will send you the one-page overview of the audit.

Related: our published studies, the AEO service, the weekly AI search newsroom, and the Lab for a check a week on your own site. Revenue Experts checks whether AI assistants name your company, and shows you the evidence. We do not promise rankings, citations or revenue; we promise the deliverable and the proof you can check.